



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/027,353	12/19/2001	Walter T. Nixon	5681-05300	9647

7590 06/20/2006

Robert C. Kowert
Conley, Rose & Tayon, P.C.
P.O. Box 398
Austin, TX 78767

EXAMINER

TSAI, SHENG JEN

ART UNIT	PAPER NUMBER
----------	--------------

2186

DATE MAILED: 06/20/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	10/027,353	NIXON ET AL.	
	Examiner	Art Unit	
	Sheng-Jen Tsai	2186	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 24 May 2006.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-38 is/are pending in the application.
- 4a) Of the above claim(s) 12 and 32 is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-11, 13-31 and 33-38 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|---|---|
| 1) <input type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. _____ |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) | 5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

1. This Office Action is taken in response to Applicants' Amendments and Remarks filed on May 24, 2006 regarding application 10/027,353 filed on December 19, 2001.

Claims 1, 2, 20-21 and 33-34 have been amended.

Claims 12 and 32 have been cancelled.

Claims 37-38 have been added.

Claims 1-11, 13-31 and 33-38 are pending for examination.

2. **Examiner's Response to Amendments and Remarks**

Applicants' amendments and remarks have been fully and carefully considered.

In response, another iteration of claim analysis based on the previously cited prior art Crater et al. (US 5,146,588) has been embarked. Refer to the corresponding sections of the claim analysis for details.

Independent claims 1, 20 and 33-34 each has been amended with additional limitations of "a memory configured to provided an addressable block operand storage space and to store block operands within the block operand storage space;" "wherein the plurality of block storage locations are configured to cache a portion of the block operand storage space of the memory and ... , such that during the first accumulation operation, the plurality of block storage locations are concurrently configured both to cache certain ones of the block operands and to accumulate the immediate result of the first accumulation operation."

Further, Applicants contend that Crater et al. (US 5,146,588) do not teach the amended limitations because Crater et al. teach the implementation of caching and

accumulation functions within separate, non-overlapping structure within cache memory (figure 3, 113) while the storage of intermediate accumulation results is performed within the redundancy accumulator (figure 4, 301), and the redundancy accumulator is not a cache of any memory.

The Examiner disagrees with this assessment for the following reason.

First, in Crater's invention, two levels of caching are provided for the memory system.

The first level of caching is provide by the memory elements (figure 3, 340~355; column 7, lines 5-20) of the cache memory unit shown in figure 1, 113, figure 2, 113 and figure 3. The memory elements (figure 3, 340~355) serves as a cache (figure 1, 113) between a mass memory entity comprising a plurality of storage devices (disk drive units, figure 1, 122-l~125-r), and a plurality of host processors (figure 1, 11~12), and all data transfers between a host processor and a redundancy group in the disk drive subsets are routed through cache accumulator memory (column 6, 7-10). In other words, cache memory (113) serves as a cache of the storage devices (figure 1, 122-l~125-r).

The second level of caching is provide by the redundancy accumulator memory (figure 4, 301), which resides inside the cache controller (figure 3, 331 and 332) and serves as a cache between a plurality of cache memory elements (figure 3, 340~355) and a plurality of interface units (figure 3, 361~364). This is the case because one of the main object of Crater et al.'s invention is to provide reliability by performing redundancy calculation of the data to form a redundancy group in the disk drive

subsets. To achieve this goal, not only that all data-transfers between a host processor and a redundancy group in the disk drive subsets are routed through cache memory 113 (column 6, 7-10), the data must also go through the redundancy accumulator memory 301 (figures 3 and 4) so that the parity can be calculated [the data elements, such as physical tracks are moved one at a time to the redundancy generator (figure 4) as they are sent via the optical file backend channels to the data storage elements in a redundancy group. This requires the use of redundancy accumulator memory (figure 4, 301) which, is used to store the intermediate result of the redundancy calculations until all of the physical tracks have been included in the redundancy calculation (column 7, lines 49-57)]. In other words, the redundancy accumulator memory (figure 4, 301) serves as a cache of cache memory elements (figure 3, 340~355).

Second, the plurality of block storage locations of the redundancy accumulator memory (figure 4, 301) are concurrently configured both to cache certain ones of the block operands and to accumulate the immediate result of the first accumulation operation [When valid prior data or prior redundancy products are stored in the addressed memory location of redundancy accumulator memory 301, these are used by redundancy generator 305 to combine with the received data byte from latch 303 in the present physical track to generate a new redundancy product which is written into redundancy accumulator memory 301 as part of the well known redundancy calculation. As each data byte is read from the presently received physical track and transmitted through latch 303 to redundancy generator 305 and multiplexer 304, this data is either directly written into the addressed memory location in redundancy

accumulator memory 301 or used to perform a redundancy calculation with the prior product stored therein and this new product is then written into the addressed memory location (column 9, lines 55-63); This requires the use of redundancy accumulator memory 301 which is used to store the intermediate result of the redundancy calculations until all of the physical tracks have been included in the redundancy calculation (column 7, lines 53-57)].

Therefore, Crater et al. clearly demonstrate and teach the amended limitations recited in these claims, and the Examiner's position regarding the status of these claims, and all those claims depending from it, remains the same as stated in the previous Office action.

Claim Rejections - 35 USC § 102

3. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

4. Claims 1-3, 5-6, 9-10, 13-23, 26, 28-31 and 33-38 are rejected under 35 U.S.C. 102(b) as being anticipated by Crater et al. (U.S. 5,146,588).

As to claim 1, Crater et al. disclose **an apparatus comprising:**

A memory [the corresponding memory is the memory elements (figure 3, 340~355; column 7, lines 5-20) of the cache memory unit shown in figure 1, 113, figure 2, 113 and figure 3] **configured to provided an addressable block operand storage space and to store block operands within the block operand storage space** [the memory

elements (figure 3, 340~355; column 7, lines 5-20) serves as a cache (figure 1, 113) between a mass memory entity comprising a plurality of storage devices (disk drive units, figure 1, 122-l~125-r), and a plurality of host processors (figure 1, 11~12), and all data transfers between a host processor and a redundancy group in the disk drive subsets are routed through cache memory (column 6, 7-10) to provide the operands of the redundancy accumulation calculations];

A functional unit [the redundancy calculator/generator, figure 4; column 7, lines 22-68; column 8, lines 1-15] **configured to perform a block operation on one or more block operands to generate a block result** [the block operation is the redundancy calculation (column 8, lines 55-68; column 9, lines 1-68; column 10, lines 1-9) that generates redundancy data for a block of data (also known as a stripe of data) distributed across N disk drives, the generated redundancy data is distributed across a block of M spared disk drives, and the N+M disk drives form a redundancy group (column 1, lines 20-23; column 2, lines 3-27; column 3, lines 3-22; column 6, lines 20-25. Data of this redundancy group constitutes a block of data, also referred to as a stripe of data in the art of disk array]; **and**

a cache accumulator memory [the corresponding cache accumulator memory is the redundancy accumulator memory (figure 4, 301), which resides inside the cache controller (figure 3, 331 and 332) and serves as a cache between a plurality of cache memory elements (figure 3, 340~355) and a plurality of interface units (figure 3, 361~364). This is the case because one of the main object of Crater et al.'s invention is to provide reliability by performing redundancy calculation of the data to form a

redundancy group in the disk drive subsets. To achieve this goal, not only that all data-transfers between a host processor and a redundancy group in the disk drive subsets are routed through cache memory 113 (column 6, 7-10), the data must also go through the redundancy accumulator memory 301 (figures 3 and 4) so that the parity can be calculated [the data elements, such as physical tracks are moved one at a time to the redundancy generator (figure 4) as they are sent via the optical file backend channels to the data storage elements in a redundancy group. This requires the use of redundancy accumulator memory (figure 4, 301) which, is used to store the intermediate result of the redundancy calculations until all of the physical tracks have been included in the redundancy calculation (column 7, lines 49-57)]. In other words, the redundancy accumulator memory (figure 4, 301) serves as a cache of cache memory elements (figure 3, 340~355)] **coupled to the memory** [figure 3 shown that the redundancy accumulator memory (figure 4, 301), which resides inside the cache controller (figure 3, 331 and 332) and serves as a cache between a plurality of cache memory elements (figure 3, 340~355) and a plurality of interface units (figure 3, 361~364), is directly coupled to the memory elements (figure 3, 340~355)] **and the functional unit** [the redundancy calculator/generator, figure 4; column 7, lines 22-68; column 8, lines 1-15], **wherein the cache accumulator memory comprises a plurality of block storage locations** [the redundancy accumulator memory is typically an n by k memory (column 8, lines 23-24)], **wherein the cache accumulator memory is configured to receive a set of one or more instructions to perform a first accumulation operation** [the read/accumulate control line shown in figure 4 indicates

the possible operations to be performed, including the accumulation operation],
**wherein a first instruction in the set uses a first address in the memory to
identify a first block operand** [the address bus shown in figure 4 identifies the block
operand];

**wherein the plurality of block storage locations are configured to cache a portion
of the block operand storage space of the memory** [the redundancy accumulator
memory (figure 4, 301), which resides inside the cache controller (figure 3, 331 and
332) and serves as a cache between a plurality of cache memory elements (figure 3,
340~355) and a plurality of interface units (figure 3, 361~364). This is the case
because one of the main object of Crater et al.'s invention is to provide reliability by
performing redundancy calculation of the data to form a redundancy group in the disk
drive subsets. To achieve this goal, not only that all data-transfers between a host
processor and a redundancy group in the disk drive subsets are routed through cache
memory 113 (column 6, 7-10), the data must also go through the redundancy
accumulator memory 301 (figures 3 and 4) so that the parity can be calculated [the
data elements, such as physical tracks are moved one at a time to the redundancy
generator (figure 4) as they are sent via the optical file backend channels to the data
storage elements in a redundancy group. This requires the use of redundancy
accumulator memory (figure 4, 301) which, is used to store the intermediate result of
the redundancy calculations until all of the physical tracks have been included in the
redundancy calculation (column 7, lines 49-57)]. In other words, the redundancy
accumulator memory (figure 4, 301) serves as a cache of cache memory elements

(figure 3, 340~355)] **and to accumulate an intermediate result of the first accumulation operation** [figure 4 shows the redundancy accumulation operations; the data elements, such as physical tracks are moved one at a time to the redundancy generator (figure 4) as they are sent via the optical file backend channels to the data storage elements in a redundancy group. This requires the use of redundancy accumulator memory (figure 4, 301) which, is used to store the intermediate result of the redundancy calculations until all of the physical tracks have been included in the redundancy calculation (column 7, lines 49-57)], **wherein the intermediate result is both a result of and an operand of the first accumulation operation** [Crater et al. teach that the operand for the redundancy calculation may come from the intermediate results stored in the redundancy accumulator memory: the multiplexers shown in figure 4, 304 & 310 select an operand from either an operand supplied from the external memory via the data input bus (the corresponding data path is DATA INPUT BUS → LATCH (303) → DATA SELECTOR (304) → multiplexer (306) → IN port of the redundancy accumulator (301)), or an operand provided by the redundancy accumulator (the corresponding data path is OUT port of the redundancy accumulator (301) → LATCH (306) → multiplexer (306) → IN port of the redundancy accumulator (301)); in the case where data is read from redundancy accumulator memory ... (column 8, lines 65-67; column 9, lines 21-36; column 9, lines 49-68; column 10, lines 1-9)] **such that during the first accumulation operation, the plurality of block storage locations are concurrently configured both to cache certain ones of the block operands and to accumulate the immediate result of the first accumulation**

operation [When valid prior data or prior redundancy products are stored in the addressed memory location of redundancy accumulator memory 301, these are used by redundancy generator 305 to combine with the received data byte from latch 303 in the present physical track to generate a new redundancy product which is written into redundancy accumulator memory 301 as part of the well known redundancy calculation. As each data byte is read from the presently received physical track and transmitted through latch 303 to redundancy generator 305 and multiplexer 304, this data is either directly written into the addressed memory location in redundancy accumulator memory 301 or used to perform a redundancy calculation with the prior product stored therein and this new product is then written into the addressed memory location (column 9, lines 55-63); This requires the use of redundancy accumulator memory 301 which is used to store the intermediate result of the redundancy calculations until all of the physical tracks have been included in the redundancy calculation (column 7, lines 53-57)].

wherein in response to receiving the first instruction in the set, the cache accumulator memory is configured to access an associativity list [the corresponding associativity list is the pointer memory, figure 4, 302, column 8, lines 16-54] **comprising an indication that a first set of the block storage locations is allocated to the first accumulation operation and** [column 8, lines 16-54; column 8, lines 55-68; column 9, lines 1-68; column 10, lines 1-9], **in response to the indication, to provide the first block operand to the functional unit from the first set of block storage locations** [column 8, lines 16-54; column 8, lines 55-68; column

Art Unit: 2186

9, lines 1-68; column 10, lines 1-9] **and to store the block result generated by the functional unit into the first set of block storage locations** [column 8, lines 16-54; column 8, lines 55-68; column 9, lines 1-68; column 10, lines 1-9].

As to claim 2, Crater et al. teach that **the cache accumulator memory comprises a dual-ported memory configured to provide the plurality of block storage locations** [figure 4 shows that the redundancy accumulator memory (301) has two ports: IN and OUT; the redundancy accumulator memory is typically an n by k memory (column 8, lines 23-24), hence providing a plurality of block storage locations].

As to claim 3, Crater et al. teach that **the cache accumulator memory [a redundancy accumulator memory, (figure 4, 301) and a pointer memory (figure 4, 302)] comprises at least two independently interfaced memory banks** [figure 3 shows that a plurality of independently interfaced memory banks (340~355) are present; column 7, lines 5-21], **wherein the cache accumulator memory is configured to provide the first block operand from a first block storage location** [corresponding to the N disk drive containing the payload data; figure 4 shows that the redundancy accumulator memory may provide an operand where the corresponding data path is OUT port of the redundancy accumulator (301) → LATCH (306) → multiplexer (306) → IN port of the redundancy accumulator (301)] **in a first one of the independently interfaced memory banks** [all data transfers between a host processor (figure 1, 11) and a redundancy group in the disk drive subsets are routed through cache accumulator memory (figure 1, 113; figure 3); a plurality of cache interfaces (figure 3, 361~364) are provided in cache accumulator memory (figure 1, 113) to interconnect

Art Unit: 2186

cache accumulator memory with host processor and the disk drive (column 7, lines 5-21); the correspondence of the memory banks and the N+M redundancy group of disk drives is established and controlled by the cluster controller (figure 1, 110-1 and 110-2; figure 2)] **and to store the block result** [the calculated redundancy data] **in a second block storage location in a second one of the independently interfaced memory banks** [the calculated redundancy data is stored back to the M spared disk drives via the corresponding cache memory banks (figure 3, 340~355) under the control of the cluster controller (figure 1, 110-1 and 110-2; figure 2)], **wherein the first set of block storage locations** [the redundancy group formed by the N+M disk drives] **comprises the first block storage location** [corresponding to the N disk drive containing the payload data] **and the second block storage location** [corresponding to the M disk drive containing the redundancy data].

As to claim 5, Crater et al. teach that **the cache accumulator memory** [a redundancy accumulator memory, (figure 4, 301) and a pointer memory (figure 4, 302)] **is configured to load a copy of the first block operand into the first set of block storage locations in the cache accumulator memory from the memory** [not only that all data-transfers between a host processor and a redundancy group in the disk drive subsets are routed through cache accumulator memory (column 6, 7-10), the data must also go through the redundancy accumulator (figures 3 and 4) so that the parity can be calculated ; a byte from a received physical track is read into latch and redundancy accumulator memory is addressed by placement of a corresponding data address on address bus ...wherein the received data on the data input is used to

Art Unit: 2186

calculate redundancy information (column 8, lines 55-67); the data elements, such as physical tracks are moved one at a time to the redundancy generator (figure 4) as they are sent via the optical file backend channels to the data storage elements in a redundancy group. This requires the use of redundancy accumulator memory (figure 4, 301) which, is used to store the intermediate result of the redundancy calculations until all of the physical tracks have been included in the redundancy calculation (column 7, lines 49-57); the "data input bus" and "address bus" shown in figure 4 further illustrate that data-transfers between a host processor and a redundancy group in the disk drive subsets must also go through the redundancy accumulator] **in response to the first block operand not being present in the cache accumulator when the first instruction is received** [... data received from the data input bus from the disk drives is used to calculate redundancy information (column 8, lines 55-68; column 9, lines 1-68; column 10, lines 1-9); Crater et al. teach that the operand for the redundancy calculation may come from within the disk drive (i.e., the external memory) [a byte from a **received physical track** is read into latch and redundancy accumulator memory is addressed by placement of a corresponding data address on address bus ...wherein the received data on the data input is used to calculate redundancy information (column 8, lines 55-67)]; .

As to claim 6, Crater et al. teach that **if all of the block storage locations in the cache accumulator memory are currently storing valid data when the first instruction is received** [When valid prior data or prior redundancy products are stored in the addressed memory location of redundancy accumulator memory 301, these are

Art Unit: 2186

used by redundancy generator 305 to combine with the received data byte from latch 303 in the present physical track to generate a new redundancy product which is written into redundancy accumulator memory 301 as part of the well known redundancy calculation. As each data byte is read from the presently received physical track and transmitted through latch 303 to redundancy generator 305 and multiplexer 304, this data is either directly written into the addressed memory location in redundancy accumulator memory 301 or used to perform a redundancy calculation with the prior product stored therein and this new product is then written into the addressed memory location (column 9, lines 55-63)], **the cache accumulator is configured to select the first set of block storage locations and to load the copy of the first block operand into the first set of block storage locations** [data copied into the redundancy accumulator memory (figure 4, 301) by cache memory controller (figure 3, 331 and 332) (column 7, lines 5-21)], **wherein the cache accumulator memory is further configured to update the indication in the associativity list** [via the pointer memory (column 8, lines 16-54)] **to indicate that the first set of the block storage locations is allocated to the first accumulation operation in response to selecting the first set of block storage locations** [(column 8, lines 55-68; column 9, lines 1-68; column 10, lines 1-9)].

As to claim 9, Crater et al. teach that **in response to loading the first block operand into the first set of block storage locations** [into the redundancy accumulator memory (figure 4, 301)], **the cache accumulator memory is configured to update a tag** [the status bit in the associated pointer memory (figure 4, 302)]

associated with the first set of block storage locations to indicate that the first block operand is stored within [the status bit in the pointer memory is set to indicate that new data is written therein (column 8, lines 16-54)].

As to claim 10, Crater et al. teach that **the cache accumulator memory** [a redundancy accumulator memory, (figure 4, 301) and a pointer memory (figure 4, 302)] **is configured to update the associativity list** [the associated pointer memory (figure 4, 302)] **in response to storing the block result generated by the functional unit, wherein the cache accumulator is configured to update the associativity list by updating a tag** [the status bit in the pointer memory] **associated with the first set of block storage locations to indicate that the block result is stored within the first set of block storage locations** [the corresponding associativity list is the pointer memory, figure 4, 302; a bit in pointer memory is set each time a byte in redundancy accumulator memory is used to store a byte of a physical track as it is transferred from cache for storage in memory elements ... (column 8, lines 16-54)].

As to claims 13-14, Crater et al. teach that **the functional unit is configured to perform a parity calculation on the block operand** [this redundancy can be a simple parity, an orthogonal parity, or a more complicated redundancy scheme such as a Reed-Solomon code (column 7, lines 38-41)].

As to claim 15, Crater et al. teach that **the functional unit is configured to calculate a parity block from a plurality of data blocks in a stripe of data, wherein the first block operand is a first one of the data blocks in the stripe of data** [processor 204-0 also causes the redundancy data across the N data disks in a

redundancy group to be generated within cache accumulator memory 113 and writes the M segments of redundancy data onto the M redundancy disks in the redundancy group (column 6, lines 20-24)).

As to claim 16, Crater et al. teach that **the functional unit is configured to perform the operation on two block-operands** [this redundancy can be a simple parity, an orthogonal parity, or a more complicated redundancy scheme such as a Reed-Solomon code (column 7, lines 38-41). Note that a parity operation needs two operands; figure 4 shows that the XOR operator (305) takes two operands].

As to claim 17, Crater et al. teach that **a first of the two block-operands** [figure 4 shows that the XOR operator (305) takes two operands] **is the first block operand stored in the cache accumulator** [from the OUT port of the redundancy accumulator memory to the latch (306) then to the XOR operator (305); "... in the case where data is read from redundancy accumulator memory ... (column 8, lines 55-68; column 9, lines 1-68; column 10, lines 1-9)] **and a second of the two block-operands is provided on a data bus coupled to provide operands to the functional unit** [figure 4, from the DATA INPUT BUS to the latch (303) and then to the XOR operator (305)].

As to claim 18, Crater et al. teach that **a first of the two block-operands** [figure 4 shows that the XOR operator (305) takes two operands] **is the first block operand stored in the cache accumulator** [from the OUT port of the redundancy accumulator memory to the latch (306) then to the XOR operator (305); "... in the case where data is read from redundancy accumulator memory ... (column 8, lines 55-68; column 9, lines 1-68; column 10, lines 1-9)] **and a second of the two block-operands is**

provided from the memory [figure 4, from the DATA INPUT BUS to the latch (303) and then to the XOR operator (305), note that the data from the memory is also routed through the DATA INPUT BUS].

As to claim 19, Crater et al. teach that **the cache accumulator memory** [a redundancy accumulator memory, (figure 4, 301) and a pointer memory (figure 4, 302)] **is configured to store a word of the block result** [prior redundancy products are stored in the redundancy accumulator memory (column 9, lines 55-68)] **during an access cycle in which cache accumulator memory also provides a word of the first block operand to the means for performing a block operation** [column 9, lines 55-68; the redundancy accumulator memory has two ports (IN and OUT) to support this function (figure 4)].

As to claim 20, Crater et al. disclose **a method of performing a block accumulation operation using a cache accumulator memory** [a redundancy accumulator memory, (figure 4, 301) and a pointer memory (figure 4, 302)] **that comprises a plurality of block storage locations** [refer to "As to claim 1"], **the method comprising:**

receiving a first command in a set of commands used to implement an accumulation operation [redundancy accumulation operation], **wherein the first command is an instruction to perform an operation on a first block operand identified by a first address in a memory** [corresponding to the address of the N disk drive containing the payload data; the address bus shown in figure 4 identifies the block operand; the corresponding memory is the memory elements (figure 3, 340~355;

Art Unit: 2186

column 7, lines 5-20) of the cache memory unit shown in figure 1, 113, figure 2, 113 and figure 3]] **and to store a result of the operation** [column 10, lines 10-41], **wherein the result is identified by a second address in the memory** [corresponding to the address of the M disk drive containing the redundancy data; column 10, lines 10-41]; **where the memory** [the corresponding memory is the memory elements (figure 3, 340~355; column 7, lines 5-20) of the cache memory unit shown in figure 1, 113, figure 2, 113 and figure 3]] **is configured to provided an addressable block operand storage space and to store block operands within the block operand storage space** [The cache memory (figure 3, 113) serves as a cache (figure 1, 113) between a mass memory entity comprising a plurality of storage devices (disk drive units, figure 1, 122-l~125-r), and a plurality of host processors (figure 1, 11~12), and all data transfers between a host processor and a redundancy group in the disk drive subsets are routed through cache memory (column 6, 7-10) to provide the operands of the redundancy accumulation calculations], **and wherein the plurality of block storage locations are configured to cache a portion of the block operand storage space of the memory** [the redundancy accumulator memory (figure 4, 301), which resides inside the cache controller (figure 3, 331 and 332) and serves as a cache between a plurality of cache memory elements (figure 3, 340~355) and a plurality of interface units (figure 3, 361~364). This is the case because one of the main object of Crater et al.'s invention is to provide reliability by performing redundancy calculation of the data to form a redundancy group in the disk drive subsets. To achieve this goal, not only that all data-transfers between a host processor and a redundancy group in the disk drive subsets

Art Unit: 2186

are routed through cache memory 113 (column 6, 7-10), the data must also go through the redundancy accumulator memory 301 (figures 3 and 4) so that the parity can be calculated [the data elements, such as physical tracks are moved one at a time to the redundancy generator (figure 4) as they are sent via the optical file backend channels to the data storage elements in a redundancy group. This requires the use of redundancy accumulator memory (figure 4, 301) which, is used to store the intermediate result of the redundancy calculations until all of the physical tracks have been included in the redundancy calculation (column 7, lines 49-57)]. In other words, the redundancy accumulator memory (figure 4, 301) serves as a cache of cache memory elements (figure 3, 340~355)].

in response to said receiving a first command:

accessing an associativity list [the corresponding associativity list is the pointer memory, figure 4, 302, column 8, lines 16-54] **comprising an indication that a first set of block storage locations is allocated to the first accumulation operation** [column 8, lines 16-54];

in response to the indication, providing the first block operand from the first set of block storage locations to a functional unit [figure 4, the redundancy generator] **and storing a block result of the operation generated by the functional unit into the first set of block storage locations** [refer to "As to claim 1"].

As to claim 21, refer to "As to claim 2."

As to claim 22, refer to "As to claim 3."

As to claim 23, refer to "As to claim 6."

As to claim 26, refer to "As to claim 9" and "As to claim 10."

As to claim 28, refer to "As to claim 13."

As to claim 29, refer to "As to claim 14."

As to claim 30, refer to "As to claim 16."

As to claim 31, refer to "As to claim 17."

As to claim 33, refer to "As to claim 1" and "As to claim 20."

As to claim 34, refer to "As to claim 1" and "As to claim 20." Further, figure 1 shows a **host computer system** (11 and 12), a **storage array** (103-1), and an **interconnect coupled to the host computer system and the storage array** [figure 1] and configured to transfer data between the host computer system and the **storage array** [column 3, lines 40-68; column 4, lines 1-68].

As to claim 35, refer to "As to claim 13" and "As to claim 17."

As to claim 36, refer to "As to claim 15."

As to claim 37, Crater et al. teach that **the apparatus as recited in claim 1, wherein the block operands stored by the memory** [the corresponding memory is the memory elements (figure 3, 340~355; column 7, lines 5-20) of the cache memory unit shown in figure 1, 113, figure 2, 113 and figure 3] **within the block operand storage space are received from one or more mass storage devices** [a mass memory entity comprising a plurality of storage devices (disk drive units, figure 1, 122-1~125-r)].

As to claim 38, refer to "As to claim 37."

Claim Rejections - 35 USC § 103

5. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

6. Claims 4, 7-8 and 24-25 are rejected under 35 U.S.C. 103(a) as being unpatentable over Crater et al. (U.S. 4,888,679), and in view of Handy, "The cache memory book: the authoritative reference on cache design," 2nd edition, Academic Press, 1998.

As to claim 4, Crater et al. do not explicitly mention that **the cache accumulator memory is configured to indicate whether a particular block operand stored in the cache accumulator is modified with respect to a copy of that particular block operand in the memory.**

However, Crater et al. do teach that "control of the cache accumulator memory is provided in control unit (figure 1, 101) by processor (figure 2, 204-0), and that the functions provided by the processor include initialization of the cache directory and other cache data structures, cache directory searching and management, cache space management, cache performance improvement algorithms as well as other cache control functions (column 6, lines 10-17).

Further, it is well known and understood in the art that certain cache management functions are required to all cache memory systems, including a mechanism to maintain the consistency between the main memory and the cache

memory, and an indicator/flag, commonly known as the “dirty bit,” is required to indicate whether the data in the cache has been modified and hence is different from the corresponding copy in the main memory.

Moreover, Handy teaches in the text book “The Cache Memory Book: The Authoritative Reference on Cache Design,” (2nd edition, Academic Press, 1998, page 66) this property of a cache memory system.

Therefore it would have been obvious for persons of ordinary skills in the art at the time of applicant’s invention to recognize the well known and understood property of the “dirty” bit that is widely adopted in cache systems, hence lacking patentable significance.

As to claim 7, Crater et al. do not explicitly mention that **the cache accumulator memory is configured to use a least recently used algorithm to select the first set of block storage locations to overwrite.**

However, Crater et al. do teach that “control of the cache accumulator memory is provided in control unit (figure 1, 101) by processor (figure 2, 204-0), and that the functions provided by the processor include initialization of the cache directory and other cache data structures, cache directory searching and management, cache space management, cache performance improvement algorithms as well as other cache control functions (column 6, lines 10-17).

Further, it is well known and understood in the art that certain cache management functions are required to all cache memory systems, including a

replacement algorithm is required in a cache system to select which entry in the cache is to be replaced when a new line is to be brought into the cache.

Moreover, Handy teaches that that the least recent used algorithm is one of the most commonly adopted scheme (page 57).

Therefore it would have been obvious for persons of ordinary skills in the art at the time of applicant's invention to recognize the well known and understood property of the least recent used algorithm that is widely adopted in cache systems, hence lacking patentable significance.

As to claim 8, Crater et al. do not explicitly mention that **if data to be overwritten in the first set of block storage locations is modified with respect to a copy of that data in the memory, the cache accumulator memory is configured to write the data back to the memory before loading the copy of the first block operand into the first set of block storage locations.**

However, Crater et al. do teach that "control of the cache accumulator memory is provided in control unit (figure 1, 101) by processor (figure 2, 204-0), and that the functions provided by the processor include initialization of the cache directory and other cache data structures, cache directory searching and management, cache space management, cache performance improvement algorithms as well as other cache control functions (column 6, lines 10-17).

Further, it is well known and understood in the art that certain cache management functions are required to all cache memory systems, including a

Art Unit: 2186

coherency policy to maintain data consistency between the main memory and the cache memory.

Moreover, Handy teaches that a write strategy is required in a cache system to deal with the situations where data is modified in either the cache or the main memory, which leads to data inconsistency between the main memory and cache. Particularly, Handy teaches that a technique, known as "write-through," in which the main memory is always updated first during all write cycles, is commonly adopted in cache system design (pages 64-65). With such a write-through policy, data consistency between the main memory and the cache will be enforced.

Therefore it would have been obvious for persons of ordinary skills in the art at the time of applicant's invention to recognize the well known and understood property of a write policy and the write-through algorithm that is widely adopted in cache systems, hence lacking patentable significance.

As to claim 24, refer to "As to claim 7."

As to claim 25, refer to "As to claim 8."

7. Claims 11 and 27 are rejected under 35 U.S.C. 103(a) as being unpatentable over Crater et al. (U.S. 4,888,679).

As to claim 11, Crater et al. teach that **the cache accumulator memory is configured to update the tag [the status bit in the pointer memory] corresponding to a second address in the memory that identifies the block result** [refer to "As to claims 9 and 10;" column 8, lines 55-68; column 9, lines 1-68; column 10, lines 1-9],

but do not explicitly mention setting the tag to equal to a first portion of address bits of a second address in the memory that identifies the block result.

However, Crater et al. teach that there is a one-to-one correspondence between a memory location in the pointer memory and each memory location in redundancy accumulator memory, and that a bit in pointer memory is set each time a byte in redundancy accumulator memory is used to store a byte of a physical track as it is transferred from cache for storage in memory elements (column 8, lines 16-54).

Further, locations in redundancy accumulator memory are also used to store intermediate results (i.e., the block results) of the redundancy calculation (column 7, lines 53-60).

Therefore it would have been obvious for persons of ordinary skills in the art at the time of applicant's invention to recognize the pointer memory essentially identifies and indicates the address of the block result, hence the claimed limitation lacks patentable significance.

As to claim 27, refer to "As to claim 11."

Conclusion

8. Claims 1-11, 13-31 and 33-38 are rejected as explained above.

9. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the

shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.


10. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Sheng-Jen Tsai whose telephone number is 571-272-4244. The examiner can normally be reached on 8:30 - 5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matthew Kim can be reached on 571-272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Sheng-Jen Tsai
Examiner
Art Unit 2186

June 8, 2006


PIERRE BATAILLE
PRIMARY EXAMINER
6/15/06